HP StorageWorks

Enterprise File Services (EFS) Clustered Gateway 3.0.1 Linux Edition

release notes

395343–002

Part number: 395343-002
First edition: August 2005

Legal and notice information

Enterprise File Services (EFS) Clustered Gateway release notes

About this document

This section describes the content reflected in this document, including:

Release notes information

These release notes contain the following sections:

Intended audience

This document is intended for customers who purchased the Linux version of HP StorageWorks Enterprise File Services (EFS) Clustered Gateway product.

Accessing future product updates

HP strongly recommends that customers sign up online using the Subscriber's choice web site: http://www.hp.com/go/e-updates.

- Subscribing to this service provides you with e-mail updates on the latest product enhancements, newest versions of drivers, and firmware documentation updates, as well as instant access to numerous other product resources.
- After signing up, you can quickly locate your products by selecting **Business support,** and then **Storage** under Product Category.

Other documentation

The following other documentation is available with this product:

- Online help/user guide (accessible via the Web browser interface)
- *HP StorageWorks Clustered File System administration guide*
- *HP StorageWorks Clustered File System setup guide*
- *HP StorageWorks Clustered File System command line reference guide*
- *HP StorageWorks Clustered File System for Linux: NFS configuration installation and administration guide*
- *HP StorageWorks Clustered File System third party software*

Additional documentation, including whitepapers and best-practices documents, is available via the HP web site: http://www.hp.com.

What is new in 3.0.1

The following new features/capabilities have been added to this release:

- Support for NFS over TCP.
- Inclusion of the QLogic 8.00.02 HBA driver (support for EVA 4000/6000/8000).
- Performance enhancements for large PSFS filesystems.
- Improvements in general I/O throughput.
- Performance enhancements for Oracle operations. The 3.0.1 release provides dramatic performance improvements to NFS exports of DB-OPTIMIZED PSFS filesystems. Improvements of 10 to 40 times greater throughput for databases as NFS clients has been measured.
- Incorporation of much of ProLiant Service Pack (PSP) 7.3 including:
  - HP Insight Manger 7.3 Agents (now pre-installed)
  - Updated BCM5700 driver (Ethernet)
  - HP fabric utilities (look in /opt/hp/hp_fibreutils)
  - The following items were not included from the PSP:
    - QLogic HBA Driver (A special QLogic 8.00.02b driver was included from another source.)
    - Intel NIC, Adaptec SCSI, CCISS Driver, VLAN
  - Smaller footprint (removal of numerous unused packages like ISDN, sound drivers, and so on.)
- New PSFS filesystem mount option to disable cluster-coherent locking. When this option is used, the filesystem provides only single node coherent locking.

  Cluster-coherent record locking can be disabled for a specific filesystem at mount time. On the Management Console, click **Advanced** on the Mount Filesystem dialog to display the Advanced Mount Options dialog. Then remove the checkmark from the Enable Cluster Coherent Record Locks option. Cluster-coherent locking will be disabled on all nodes in the cluster.

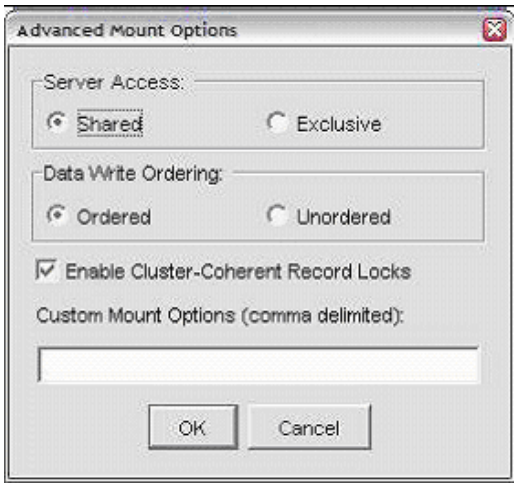

NOTE: The Shared/Exclusive and Ordered/Unordered options are now on the Advanced Mount Options dialog. They were previously on the Mount Filesystem dialog.

Cluster-coherent locking can also be disabled when mounting a filesystem from the command line by specifying the `localbrlocks` option.

NOTE: This mount option should be used with extreme caution. When cluster-coherent record lock operations are disabled, the record locks established against a file in the filesystem by each node in the cluster are invisible to all other nodes in the same cluster. Before using this option, ensure that either all processes performing record lock operations in the affected filesystem are running on the same single node in the cluster, or that the necessary mutual exclusion semantics for record locking have been implemented in some other way.

Upgrading to 3.0.1

Upgrades to HP Clustered Gateway 3.0.1 are supported only from the HP Clustered Gateway 3.0.0 releases.

Back up the existing cluster

Before starting the upgrade, HP strongly recommends that you back up your existing cluster:

- Back up */var/hpcfs* and */etc/hpcfs* on each server. It is important to preserve the original cluster configuration in case you need to back out of the upgrade.
- Back up all PSFS filesystems for disaster recovery purposes.

Upgrade or re-image

There are two ways to upgrade existing 3.0.0 clusters to 3.0.1:

- Re-imaging—With this procedure, each node in the cluster is re-imaged using a 3.0.1 Quick Restore DVD. This process ensures that you have the latest Clustered Gateway software and management tools. This process requires some extra effort because you must restore the state of the node (hostname, IP addresses, etc.) after the node is re-imaged.

NOTE: Using the save/restore system state tool in the Clustered Gateway Control Panel is not supported for this task, because it restores a significant amount of data that conflicts (corrupts) files in the new image.

- Clustered File System Upgrade—This option allows you to upgrade the HP Clustered File System software without re-imaging the nodes. This process does not yield the same results as re-imaging. In particular, none of the items from the ProLiant Service Pack (PSP) 7.3 are included. The PSP can be installed after the upgrade by downloading it from the HP web site. If this approach is taken, ensure the QLogic Driver from the PSP is *not* installed.

Note for upgrades on systems using certain storage arrays

This note applies to the following storage arrays:

- HP XP arrays
- HP EVA and MSA arrays
- EMC Symmetrix

The HP Clustered Gateway 3.0.1 release fixes defect 9364, in which HP Clustered Gateway incorrectly interpreted the Control Unit (CU) and LUN reported by the HP XP array. Because this information is used to create the disk UIDs that identify each LUN, the disk UIDs may change in the 3.0.1 release.

To resolve this situation, you will need to re-create your membership partitions using the new disk UIDs and validate that the PSFS filesystems are mounted in the correct locations.

Upgrade HP Clustered Gateway and retain the membership partitions

This procedure is a rolling upgrade and does not require any cluster downtime. Each server is removed from the cluster, upgraded, and then returned to the cluster.

Rolling upgrade considerations

Consider the following when upgrading HP Clustered Gateway to the 3.0.1 release:

- Servers can be upgraded one-at-a-time during the rolling upgrade. We recommend that you upgrade the server with the lowest IP address first. Then continue to upgrade the servers in ascending order of IP address.
- During the rolling upgrade, the cluster will include a mix of servers running 3.0.1 and the earlier release (3.0.0). As long as the cluster includes at least one active server running 3.0.0, other servers running 3.0.0 will be able to join the cluster (for example, following a reboot). However, after the last server running 3.0.0 has left the cluster, only servers running 3.0.1 will be able to join the cluster.

  Although HP Clustered Gateway supports a mix of servers running 3.0.1 and 3.0.0 during the rolling upgrade, on-going operation of a mixed-version cluster is not supported. All servers should be upgraded to 3.0.1 as quickly as possible.
- For best results while the cluster includes a mix of 3.0.1 and 3.0.0 servers, all UI operations (either via the Management Console or the **mx** command-line utility) should be connected to a 3.0.0 server.
- You do not need to run **mxconfig** when upgrading to 3.0.1. If you want to change any of the parameters that you originally specified with **mxconfig**, first upgrade all servers to 3.0.1. After the upgrade is complete, stop HP Clustered Gateway on each server (use **/etc/init.d/pmxs stop**), and then rerun **mxconfig**.

Upgrade from HP Clustered Gateway 3.0.0 and retain membership partitions

Complete the following steps to upgrade to HP Clustered Gateway 3.0.1. Repeat this procedure on each server in the cluster.

NOTE: Upgrade the server with the lowest IP address first. Then continue to upgrade the servers in ascending order of IP address.

---

1. Download the 3.0.1 upgrade package from the web.
2. Untar the package using: tar -zxf cgw301upgrade.tar.gz.
3. Type `cd hpcg`.
4. Type `./upgrade`.

NOTE: The System will reboot once during the upgrade process.

---

5. After the reboot, you can remove the hpcg directory.

All servers in the cluster are now operational; however, you need to upgrade your PSFS filesystems to the format used in HP Clustered Gateway 3.0.1. See the section, "Upgrade 3.0.0 PSFS filesystems" on page 8 for more information.

Upgrade HP Clustered Gateway and re-create the membership partitions

This procedure upgrades HP Clustered Gateway to the 3.0.1 release and also re-creates the membership partitions to resolve situations in which the partitions were created with incorrect UIDs under earlier versions of HP Clustered Gateway.

Upgrade from HP Clustered Gateway 3.0.0 and re-create membership partitions

This procedure describes how to upgrade to HP Clustered Gateway 3.0.1 and re-create your membership partitions. To perform the upgrade, complete the following steps:

1. While HP Clustered Gateway is running, ensure that each PSFS filesystem has a label. See the HP Clustered Gateway Administration Guide for information about labeling filesystems.
2. Run the following command to obtain a list of the disks currently imported into the cluster and print the output. You will need to re-import these disks after the update to 3.0.1.

   ```
   /opt/hpcfs/bin/sandiskinfo -ial
   ```
3. Locate each PSFS filesystem on the output from sandiskinfo. Record the label and mountpoint on which the filesystem is mounted. You can obtain this information by running the `mx fs status` command.
4. Run the following command to obtain a list of any disks that are not currently imported into the cluster and print the output. These disks should not be included when you re-import disks after the upgrade is complete.

   ```
   /opt/hpcfs/bin/sandiskinfo -ual
   ```
5. Either re-image or upgrade each of the nodes in the cluster. When the nodes are rebooted and try to form a cluster, they should fail to do so because the membership partition UID's have changed.
6. Run the following command to obtain a list of unimported disks to see the new UID's.

   ```
   /opt/hpcfs/bin/sandiskinfo -ual
   ```
7. Run the `mxconfig` utility on one server. Skip the initial screens and go to the Membership Partition Setup window.
8. Re-create the membership partitions:
   a. Remove the existing membership partitions. Consult the sandiskinfo output obtained in step 11 to locate the partitions. (Look for the PSMP/Active partitions.)
   b. Add back the membership partitions that you removed in step a.

      (Note that the partitions may now have different UIDs.)
   c. When you have completed your selections, click **Done** on the Membership Partition Setup window.
9. Export the configuration to all servers in the cluster.
10. Start HP Clustered Gateway on each server. Run the following command:

    ```
    # /etc/init.d/pmxs start
    ```
11. Using the sandiskinfo output that you obtained in step 6, locate the disks containing PSFS filesystems. These disks need to be re-imported into the cluster. (Compare the disks on your list with the sandiskinfo -ual output that you saved in step 4, before the upgrade. Disks that were not imported prior to the upgrade should not be imported now, even if they contain PSFS filesystems.)
12. Connect to the Management Console from a server on which you started HP Clustered Gateway and re-import the disks containing PSFS filesystems into the cluster.
13. Restart HP Clustered Gateway on each server to cause the PSFS filesystems to be mounted.

    ```
    # /etc/init.d/pmxs restart
    ```
14. Verify that all PSFS filesystems are mounted on the correct mount point on the appropriate servers. To see the currently assigned paths, run the following command:

    ```
    mx fs status
    ```
15. Compare the current assignments with the sandiskinfo -ial output that you saved before the upgrade. If necessary, reassign mount points so that each filesystem is mounted as it was before the upgrade to 3.0.1. You can use either of these methods to reassign mount points:
    - On the Management Console, right-click the filesystem, select **Mount**, and then specify the correct mountpoint.
    - From the command line, use either of the following commands:

      ```
      mount -t psfs -o <options> /dev/psd/devicepYY <mountpoint>
      mx fs mount [--persist] [--activate] [--options <option>, <option>,...] --path
      <path> <filesystem> <server> ...
      ```

All servers in the cluster are now operational; however, you need to upgrade your PSFS filesystems to the format used in HP Clustered Gateway 3.0.1. See the next section, "Upgrade 3.0.0 PSFS filesystems" for more information.

Upgrade 3.0.0 PSFS filesystems

HP Clustered Gateway 3.0.1 provides a new on-disk filesystem format. This format includes the full-zone bitmap feature, which enhances performance on large filesystems.

PSFS filesystems created on HP Clustered Gateway 3.0.1 use the new on-disk format and are not compatible with HP Clustered Gateway 3.0.0 or earlier releases.

PSFS filesystems created on HP Clustered Gateway 3.0.0 or earlier releases must be upgraded to the new format after the 3.0.1 release has been installed on all servers in the cluster.

First unmount all PSFS filesystems and then run the following command:

```
psfsck --upgrade-fs <device>
```

Issues resolved in 3.0.1

This release includes the following new fixes:

- Defect QXCR1000280201. EVA 4000/6000/8000 array support

  The QLogic driver included with the product does not support the EVA 4000/6000/8000 arrays. A new driver must be obtained from HP. The required driver is unique to the HP Clustered Gateway product.

- Defect QXCR1000280207. Missing qlogic HBA config in fc_pcitable

  Before configuring the HP Clustered File System software, run the following command:

  ```
  /opt/hpcfs/lib/chhbadriver qla2xxx-8.00.00
  ```

  This command results in two new lines related to the QLogic driver being added to the end of this file:

  ```
  /etc/hpcfs/fc_pcitable
  ```

  If you are using the QLogic driver in failover mode, you must edit this file and change the text in the second to last line of the file:

  From

  ```
  ql2xfailover=0
  ```

  To

  ```
  ql2xfailover=1
  ```

- Defect 8289. After a kernel panic on one node, HP Clustered File System could keep running on that node; the node would not be fenced, as it should be.

- Defect 8370. The PSFS filesystem took too long to allocate a block on large, relatively full filesystems. To resolve this problem, the filesystem now includes a Full Zone BitMap feature that reduces the amount of data the filesystem needs to read when allocating a block. If you are upgrading to HP Clustered File System 3.0.1, you will need to upgrade existing PSFS filesystems to use the Full Zone BitMap feature.

- Defect 8888. File unlink operations were slow.

- Defect 9035. The `mxmpconf` utility includes repair actions, such as Remove and Resilver, that present a list of current membership partitions for selection. In the previous release, the selection menu was incorrect, preventing the user from selecting the appropriate membership partition.

- Defect 9067. The Management web server was updated for a security fix.

- Defect 9095. The system hung after the QLogic driver made several failover attempts.

- Defect 9172. The PSFS filesystem failed to release a lock, causing a race condition that could result in either runtime deadlocks or hangs when nodes were shut down.

- Defect 9189. A NullPointerException occurred on the Management Console when the option to configure the cluster was selected on the Connection dialog.

- Defect 9201. A flaw existed in the mxreg libraries used by vstatd in which shutdown would sometimes cause threads to be reaped more than once, causing memory corruption leading to a SEGV.

- Defect 9212. Bitmap data structures could consume all kernel vmalloc space. The size of the data structures was reduced in the 3.0.1 release to allow more storage to be attached.

- Defect 9216. After a node was power-cycled, it experienced a disk import error and was unable to access certain SAN disks.

- Defect 9252. A flaw in an internal library component could cause mxregd to hang in some rare circumstances.

- Defect 9256. Evicted filesystems on psv devices were not always fenced correctly.

- Defect 9287. When a mount option was specified for a PSFS filesystem, the default values for the other mount options were not always set correctly.

- Defect 9295. A mount option was needed to disable cluster-coherent record locking.

- Defect 9298. A write to filesystem mounted with the DBOPTIMIZE option would wait for other writes to the same file to finish.

- Defect 9324. An error message was needed for attempts to mount a PSFS filesystem on an earlier version of HP Clustered File System that did not support the filesystem's version or features.

- Defect 9342. The mx utility supplied with HP Clustered File System 3.0.0 could not send commands to servers running HP Clustered File System 2.7.

- Defect 9364. On HP XP arrays, HP Clustered File System incorrectly interpreted the Control Unit (CU) and LUN reported by the array. This information is used to create disk UIDs to identify each LUN. When you upgrade to HP Clustered File System 3.0.1, you need to re-create your membership partitions as described in the *HP StorageWorks Clustered File System Setup Guide*.

- Defect 9366. The Info column on the Create Filesystem window displayed the word "Basic" instead of showing the disk UIDs.

- Defect 9396. The mx utility did not properly recognize certain command-line options when running on HP Clustered File System 2.7.x or earlier.

- Defect 9411. The Disk Deport window could not be resized.

Open issues

Open issues are divided into the following sections:

HP Clustered Gateway general

No MPIO software is turned on by default [QXCR1000280209]

It is likely that multiple paths to LUN's exist. You need to enable some kind of MPIO software. The two choices are:

- The built-in mxmpio software. This software currently works with arrays with active:active configurations (for example, HP's XP array). The new EVA's have a form of active:active that is not fully supported by the mxmpio software. Namely, the performance versus non-performance paths are not correctly distinguished. To enable this solution, refer to the administration manual for the product.
- The built-in QLogic failover driver. To enable this solution:
  1. Edit the file `/etc/hpcfs/fc_pcitable`
  2. Find the line near the end of the file that has this text in the middle of the line:

     `ql2xfailover=0`
  3. Change the above string to this:

     `ql2xfailover=1`

Windows may not launch from CGW Control Panel during initial setup [QXCR1000280204]

When using the Clustered Gateway Control panel from the desktop it is possible that some of the windows/screens do not appear. This is likely because the host name of the system was changed. The X-Windows system was started with one host name and now that the host name has changed it may not allow windows from the "new" host to be displayed. To change the host name, log off and log back on. Alternatively, window security may be disabled by issuing this shell command: `xhost +`.

Initial Startup requires correct replies to detection of new hardware [QXCR1000280199]

During initial startup, the administrator might receive a series of screens stating that new hardware has been discovered and asks if the hardware should be configured. The administrator should select NO for QLogic and NIC configuration questions. The administrator should select YES for all other hardware configuration questions.

Performance problem with integrated EVA snapshot (Business Copy) [QXCR1000280203]

When an HP EVA Business Copy Snapshot is initiated from the Clustered Gateway GUI or command line it could take a long time, or even timeout and fail. This can happen when there are a large number of VDisk's managed by the CommandView EVA software. The total VDisk count is impacted by all arrays managed by CommandView EVA, not just the array for which the snapshot applies. There is no known workaround for this problem. Check the HP web site for a patch to address this issue.

HP Clustered File System 3.0.1

Service monitor attempts to start before filesystem is mounted [982]

PSFS filesystems can be configured to be automatically mounted when the system is booted. This configuration is called a *persistent mount*. When a server is booted, HP Clustered File System does not wait for the persistent mounts to complete before it begins other cluster operations.

This behavior can cause a virtual host to move among servers under the following circumstance:
A service monitor is configured for an application that depends on a persistently-mounted filesystem, and the application is running on all servers.

If all servers are booted at once, the service monitor might attempt to run on the primary server before the filesystem is mounted. When the service monitor fails, the virtual host associated with the monitor might move to another server.

This behavior is repeated until the filesystem is mounted somewhere in the cluster. The actual behavior of the monitor and virtual host are dependent on the application, and on the activities specified in the scripts associated with the monitor.

To prevent this situation, create a SHARED_FILESYSTEM device monitor on the filesystem mount. This device monitor prevents the virtual host from being established on any server that has not completed the filesystem mount.

Resized log does not display properly [1615]

When you use the Set Log Length option to set the number of lines to appear on the Server Log window, the Management Console does not refresh the display. Close the Server Log window, and then reopen it to see the resized log.

Management Console does not display persistent mounts [1717]

Certain error conditions can cause a server to lose access to a PSFS filesystem. When you unmount the evicted filesystem, the Management Console no longer displays persistent mount information for that filesystem. Log out of the Management Console and then log back in to see the information.

HP Clustered File System does not shut down applications [1735]

When HP Clustered File System is shut down, it terminates processes that have open files on PSFS filesystems. To avoid this problem, we recommend that you stop applications that are using the filesystem before you shut down HP Clustered File System.

Import operation fails for more than 63 LUNs [2001]

If you attempt to import a large number of LUNs in a single operation, the import succeeds only for the first 63 LUNs. Repeat the operation to import the remaining LUNs. Be sure to limit each operation to 63 or fewer LUNs.

Password prompt requires old password [7823]

When you change the admin password on the Management Console and then export it to other servers in the cluster, you are prompted for the password on those servers. Specify the old password, not the newly assigned password. (The old password is on the servers until the export operation is complete.)

Application tab shows unknown state as inactive [7838]

The Application tab on the Management Console uses the same unavailable pause icon for monitors that are in either the unknown state or the inactive state.

Workaround: To determine the actual state of the monitor, check the monitor on the Servers or Virtual Hosts tabs.

Mxcheck asks for FibreChannel switch information [8473]

When you run mxcheck, it asks for the names or addresses of the FibreChannel switches in the cluster. It uses this information to test the access to the switches. If are not placing FibreChannel switches under cluster control, or you do not want to test switch access at this time, press Enter at the prompt. Mxcheck continues to execute.

Volume Properties window does not show snapshot properties [8477]

The Volume Properties window shows properties for a basic or dynamic volume as well as the filesystem mounted on the volume, but does not show properties for a snapshot mounted on the volume.

To see information about a snapshot, use the utilities provided with the web management appliance.

Relocated virtual host is not available immediately [9021]

When the primary server is changed for a virtual host that is not associated with a monitor, the Management Console might report that the virtual host is DOWN. Within 15 minutes, the new server notifies the HP Clustered File System that it is the primary for the virtual host and the state of the virtual host changes to UP.

If a monitor is associated with the virtual host, the monitor immediately notifies HP Clustered File System when the primary server is changed and the state of the virtual host remains UP.

Configuration can be reset when monitor is updated [9061]

If the –type attribute is included in an mx device update or mx service update command, any monitor attributes not specified on the command line reset to the default values.

Psd device is not visible on a node [9100]

Under certain conditions, a psd device might not be visible on a server in the cluster; however, the other servers can view the device. This can occur due to transient conditions that prevented the device from being bound on that server. If so, then the following workaround might recover the device.

Workaround: Run the following command as root on the server where the device is not visible:

```
/opt/hpcfs/tools/diskupdate
```

NFS issues

NFS configuration service might be terminated [8620]

It might be possible for the NFS configuration service to be terminated mid-configuration action due to server failure or another non-recoverable error. If this occurs, it might be possible for export group entities to exist in the cluster configuration, but to not have the corresponding export records.

To correct this situation, it is necessary to re-synchronize the cluster configuration with the configuration stored in the NFS configuration service's data store. Often, deleting the export group entities and starting over is sufficient. If this situation occurs and you need assistance, contact HP Technical Support.

Running NFS client and its Virtual NFS Service on the same node [8740]

When an NFS client is running on a node that is also serving as the primary for the Virtual NFS Service used by that client, the client might quietly lose its locks when a failover occurs. To avoid this situation, do not run NFS clients on nodes that are configured as the Virtual NFS Service used by the clients.

Clients might experience delays [8799]

NFS client implementations based on Linux 2.4, Linux 2.6.5, Linux 2.6.9, and potentially other operating systems and releases, might incur the delay of an RPC timeout (on the order of 30 seconds) when granted a previously blocked lock.

This problem is caused by a client-side bug that causes the servers message granting the lock to be rejected due to authentication failure. The CFS-Linux server promptly sends the grant message and is not responsible for the potential 30-second delay.

SLES9 Linux NFS client might pause uninterruptibly [8808]

NFS client applications that use blocking file locks (via fcntl()) sometimes pauses uninterruptibly for 30 seconds and then (if they are not killed in the meantime) resume running normally. This problem is caused by a bug in the SLES9 (2.6.5 kernel) Linux NFS client.

Workaround: HP has provided a version of the SLES9 NFS client with this problem fixed. It is licensed under the GPL and is available from HP Technical Support.

Clients might fail a blocked lock request [8813]

NFS client implementations based on Linux 2.4, Linux 2.6, and potentially other operating systems and releases, might fail a blocked lock request when the server recovers after a failover. The lock failure is exposed by the NFS client to applications. A more robust client implementation would reissue the pending lock request after reclaiming any previously granted locks.

NFS operations can time out [8833]

On very large or busy clusters (busy in this context means a cluster with many virtual hosts), it is sometimes possible to experience a timeout when attempting an operation on a virtual host or an export group. The error messages read:

```
"Error: Command timed out without a response from clusterpulse."
```

Timeouts are most likely to happen when you are specifying a number of commands from the mx command line (for example, running *mx exportgroup delete group7 group8 group9* on a very busy cluster or issuing a large number of commands sequentially in a script).

If you see this error, it is safe to retry the failed commands again. There might be warnings about invalid objects (vhosts or Export Groups), but these can be ignored.

If this error persists on a smaller cluster (one with eight or fewer servers), or on a cluster that is not busy, contact HP Technical Support for help in determining why the error occurs.

Probeseverity value must be lower case [9041]

The value of the probeseverity option to the exportgroup command must be specified in lower-case letters. The command does not report an error if upper-case letters are typed; however, the option is set to the default value.

File locks cannot be acquired when NFS is not running [9081]

The NSM (Network Status Monitor) support is provided by /opt/hpcfs/sbin/vstatd, which replaces the traditional /usr/sbin/rpc.statd. vstatd is cluster-aware and requires membership in an active cluster to run. File locks cannot be acquired over NFS while NSM support is unavailable; lock requests are rejected with ENOLCK.

Export Group Details dialog does not have a scrollbar [9147]

When you click the Edit button to edit an export record on the Export Group window, the Export Record Details dialog lists the client names assigned to the record. Because the Client Names window does not have a scrollbar, you can access only the first 23 clients in the list.

Workaround: If you need to edit a client that is not displayed in the Client Names window, make your changes directly on the export record, which is shown on the NFS Exports tab of the Export Group window.

Operating system and environment issues

Problem with running the Management Console remotely [307]

If you run the Management Console with your display set to a remote X host, the Management Console might not display properly the first time you run it. Exit the Management Console, and then reset it to correct the display.

Instead of accessing the Management Console in this manner, we recommend that you install the Management Console and run it locally. The Management Console is available for both Linux and Windows systems.

Parent dentries are not revalidated [455]

During path lookups, the operating system does not revalidate parent dentries when they are specified as ".." components. The operating system assumes that the parent dentry associated with ".." is always valid; however, this is not always the case in a distributed environment.

For example, server A might have a process that executes cd /a/b/c. Server B has a process that renames /a/b/c to a/c and removes directory b. If the process on server A executes cd .., it is allowed to change directory into directory b, which no longer exists. Any operations that attempt to modify the contents of directory b fail with the appropriate error.

To get out of directory b, either execute another cd .. command or change the directory using an absolute pathname.

Problem with unmounting filesystems on the same mountpoint [701]

Linux allows two filesystems to be mounted on the same mountpoint. The second filesystem to be mounted hides the first filesystem. If you attempt to unmount the first filesystem, the second filesystem becomes unmounted because the first filesystem is not visible to the operating system. This problem can occur with any type of filesystem.

To avoid the problem, do not mount two filesystems on the same mountpoint, or unmount the second filesystem, and then unmount the first filesystem.

Mount and unmount might fail if run in parallel [1041]

If you run several mount or umount commands in parallel, an operation might fail with the following error message:

```
Cannot create link /etc/mtab~
Perhaps there is a stale lock file?
```

The mount or umount operation succeeds; however, the /etc/mtab file is not updated.
Run the command again to update the file.

Kernel oops can occur when importing many disks [1843]

If a large number of disks are imported in parallel on a system using QLogic HBAs, a kernel oops might occur. This problem is caused by a bug in the QLogic driver. If this problem occurs, reboot the affected server.

QLogic HBA driver can disable HBAs [5024]

When the QLA driver is unable to log into the FC switch for a long period of time (about 50 minutes), the driver can disable the HBA. When the switch is restored, the HBA remains disabled. Reboot the node to re-enable the HBA.

QLogic driver can exhaust kernel memory [5814]

Under certain high-stress I/O conditions, the QLogic FibreChannel driver can exhaust kernel memory.

This condition typically occurs while a server is running at very high I/O rates, and is possibly triggered by storage subsystem error recovery. During this period the system administrator might see the messages that follow:

```
Apr 30 10:06:47 qar54s1 kernel: scsi(1:5:12): QUEUE FULL status detected
0x1c-0x928, pid=128300970.
Apr 30 10:06:47 qar54s1 kernel: scsi(1:5:4): QUEUE FULL status detected
0x1c-0x928, pid=128300892.
```

Kernel memory exhaustion can result in the failure of applications or ejection of the server from the cluster. If this condition is encountered, we recommended that the command tag queue depth be reduced from the default of 32 to 16 or less.

QLogic Switch login can cause HP Clustered File System failures [7435]

HP Clustered File System does not start if you are logged into a QLogic FibreChannel switch and *admin start* is set. A message, such as the following appears:

```
Switch address is not responding to SNMP SET requests.
Verify the configured community string has SNMP write privileges.
The FibreChannel switches are not responding to SNMP requests.
Verify that the switches are configured to accept SNMP access from this server
using the configured community string. The FibreChannel switches and HP Clustered
File System must be configured with the same community string.
```

If you log into the switch with *admin start* while HP Clustered File System is running, this switch problem can affect fencing operations and might cause the cluster to hang.

Workaround: Avoid putting the switch into *admin start* mode while HP Clustered File System is running. If it is necessary to use this mode while HP Clustered File System is running, limit the use of the mode as much as possible.

Problem with QLogic 8.00.01 driver [8205]

HP does not support the use of QLogic 8.00.01 driver with HP Clustered File System because of problems the driver encounters with SCSI Inquiry commands. Instead, HP requires that you use the default 8.00.00 driver provided with HP Clustered File System.

Partition tables can be constructed incorrectly [8346]

When partitioning a disk, ensure that the first partition begins at an offset beyond sector one on the disk. This is necessary because the Linux kernel supports a number of different partition table formats, and some of them make use of sector one in addition to sector zero. If sector one is a data sector in a partition, it is possible that user data written to that sector might subsequently cause the Linux kernel to recognize a different partition table format than is actually there, which could result in the partitions being constructed incorrectly by the Linux kernel.

HP Clustered File System 3.0.1 and heterogeneous storage

The HP Clustered File System supports non-HP storage. The following open issues might be helpful when diagnosing problems with non-HP storage:

Failover can take too long with IBM RDAC software [7250]

When the IBM RDAC software is being used, the time required for failover is dependent on the RDAC software. If failover times are too long, contact HP support for assistance in tuning suggestions.

Contacting HP

Please contact us if you encounter any installation, configuration, or usage issues with HP Clustered Gateway, or if you have any comments or suggestions. Contact information and the latest product information and documentation are available on the HP Web site:

http://www.hp.com